OFFICE OF THE DIRECTOR OF NATIONAL INTELLIGENCE

Machine Intelligence from Cortical Networks (MICrONS) Proposers' Day Conference

LEADING INTELLIGENCE INTEGRATION

Office of Safe and Secure Operations
Intelligence Advanced Research Projects Activity

Proposers' Day Agenda – Morning

8:00AM – 8:30AM	Registration	
8:30AM – 8:35AM	**Welcome and opening remarks**	R. Jacob Vogelstein Program Manager
8:35AM – 9:00AM	IARPA Overview	Peter Highnam Director, IARPA
9:00AM – 10:00AM	MICrONS Program Overview	R. Jacob Vogelstein Program Manager
10:00AM – 10:30AM	Break	
10:30AM – 11:00AM	Doing Business with IARPA	Tarek Abboushi IARPA Acquisitions
11:00AM – 12:00PM	MICrONS Program Feedback and Q&A	R. Jacob Vogelstein Program Manager
12:00PM – 5:00PM	Offerors' Capabilities Briefings and Posters	Attendees (No Government)

MICrONS Team

- Programmatic SETA: Jimmy Baker
  - Retired Air Force
  - 20 years experience in IC
- Technical SETA: David Markowitz
  - PhD in molecular biology and neuroscience from Princeton University
    - Advised by David Tank
    - Collaborated with John Hopfield and Carlos Brody
  - Postdoc in primate neurophysiology at NYU
    - Advised by Bijan Pesaran
    - Studied computation in prefrontal cortex

Proposers’ Day Agenda – Afternoon

12:00PM – 1:30PM	Lunch / Poster Session	Attendees (No Government)
1:30PM – 4:00PM	Presentation Session	Attendees (No Government)
1:30PM – 2:00PM	Technical Area 1	
2:00PM – 2:30PM	Technical Area 2	
2:30PM – 3:00PM	Break	
3:00PM – 3:30PM	Technical Area 3	
3:30PM – 4:00PM	Technical Area 4	
4:00PM – 5:00PM	Poster Session	Attendees (No Government)

Today's Goals

1. Familiarize participants with IARPA and with the MICrONS program concept
2. Solicit feedback and questions (more about this on the next slide)
3. Foster networking and discussion of synergistic opportunities and capabilities among potential program participants (AKA "teaming")

Asking and Answering Questions

- Today
  - Please ask questions and make suggestions: this is your chance to influence the design of the program
  - We appreciate and seek constructive feedback on any/all aspects of the program design and program metrics
  - Record your questions and comments on the note cards provided and submit them to IARPA staff during the break
- Tomorrow (and beyond)
  - Questions will only be answered in writing on the program website
  - Once a BAA is released, questions can only be submitted to the email address provided in the BAA

Disclaimer

- These presentations are provided solely for information and planning purposes
- The Proposers' Day Conference does not constitute a formal solicitation for proposals or abstracts
- Nothing said at Proposers' Day changes the requirements set forth in a BAA
- A BAA supersedes anything presented or said by IARPA at the Proposers' Day

Proposers' Day Agenda

8:00AM – 8:30AM	Registration	
8:30AM – 8:35AM	Welcome and opening remarks	R. Jacob Vogelstein Program Manager
8:35AM – 9:00AM	**IARPA Overview**	**Peter Highnam Director, IARPA**
9:00AM – 10:00AM	MICrONS Program Overview	R. Jacob Vogelstein Program Manager
10:00AM – 10:30AM	Break	
10:30AM – 11:00AM	Doing Business with IARPA	Tarek Abboushi IARPA Acquisitions
11:00AM – 12:00PM	MICrONS Program Feedback and Q&A	R. Jacob Vogelstein Program Manager
12:00PM – 1:30PM	Lunch / Poster Session	
1:30PM – 4:00PM	Proposers' Capabilities Briefings	Attendees (No Government)
4:00PM – 5:00PM	Poster Session	Attendees (No Government)

OFFICE OF THE DIRECTOR OF NATIONAL INTELLIGENCE

Intelligence ARPA (IARPA) Overview Briefing

LEADING INTELLIGENCE INTEGRATION

Last updated May 2014

Office of the Director of National Intelligence

Central Intelligence Agency

Defense Intelligence Agency

Department of State

National Security Agency

Department of Energy

National Geospatial-Intelligence Agency

Department of the Treasury

National Reconnaissance Office

Drug Enforcement Administration

Army

Federal Bureau of Investigation

Navy

Department of Homeland Security

Air Force

Coast Guard

Marine Corps

IARPA Mission and Method

IARPA's mission is to invest in high-risk/high-payoff research that has the potential to provide the U.S. with an overwhelming intelligence advantage over our future adversaries

- **Bring the best minds to bear on our problems**
  - Full and open competition to the greatest possible extent
  - World-class, rotational, Program Managers
- **Define and execute research programs that:**
  - Have goals that are clear, measureable, ambitious and credible
  - Employ independent and rigorous Test & Evaluation
  - Involve IC partners from inception to finish
  - Run from three to five years

Office of Incisive Analysis

"Maximizing Insight from the Information We Collect, in a Timely Fashion"

Large Data Volumes and Varieties

Providing powerful new sources of information from massive, noisy data that currently overwhelm analysts.

Social-Cultural and Linguistic Factors

Analyzing language and speech to produce insights into groups and organizations.

Improving Analytic Processes

Dramatic enhancements to the analytic process at the individual and group level.

Office of Smart Collection

"Dramatically Improve the Value of Collected Data"

Novel Access

Provide technologies for reaching hard targets in denied areas

Asset Validation and Identity Intelligence

Detect the trustworthiness of others

Advance biometrics in real-world conditions

Tracking and Locating

Accurately locate HF emitters and low-power, moving emitters with a factor of ten improvement in geolocation accuracy

Office of Safe and Secure Operations

"Counter Emerging Adversary Potential to Deny our Ability to Operate Effectively in a Globally-Interdependent and Networked Environment"

Computational Power

Revolutionary advances in science and engineering to solve problems intractable with today's computers

Trustworthy Components

Getting the benefits of leading-edge hardware and software without compromising security

Safe and Secure Systems

Safeguarding mission integrity in a hostile world

How to engage with IARPA

- **Website: www.IARPA.gov**
  - Reach out to us, especially the IARPA PMs. Contact information on the website.
  - Schedule a visit if you are in the DC area or invite us to visit you.

- **Opportunities to Engage**:
  - **Research Programs**
    - Multi-year research funding opportunities on specific topics
    - Proposers' Days are a great opportunity to learn what is coming, and to influence the program
  - **"Seedlings"**
    - Allow you to contact us with your research ideas at any time
    - Funding is typically 9-12 months; IARPA funds to see whether a research program is warranted
    - IARPA periodically updates the topics of interest
  - **Requests for Information (RFIs) and Workshops**
    - Often lead to new research programs, opportunities for you to provide input while IARPA is planning new programs

Concluding Thoughts

- **Our problems are complex and truly multidisciplinary**
- **Technical excellence & technical truth**
  - Scientific Method
  - Peer/independent review
  - Full and open competition
- **We are always looking for outstanding PMs**
- **How to find out more about IARPA**:

  www.IARPA.gov
- **Contact Information**

  Phone: 301-851-7500

Proposers' Day Agenda

8:00AM – 8:30AM	Registration	
8:30AM – 8:35AM	Welcome and opening remarks	R. Jacob Vogelstein Program Manager
8:35AM – 9:00AM	IARPA Overview	Peter Highnam Director, IARPA
9:00AM – 10:00AM	**MICrONS Program Overview**	**R. Jacob Vogelstein Program Manager**
10:00AM – 10:30AM	Break	
10:30AM – 11:00AM	Doing Business with IARPA	Tarek Abboushi IARPA Acquisitions
11:00AM – 12:00PM	MICrONS Program Feedback and Q&A	R. Jacob Vogelstein Program Manager
12:00PM – 1:30PM	Lunch / Poster Session	
1:30PM – 4:00PM	Proposers' Capabilities Briefings	Attendees (No Government)
4:00PM – 5:00PM	Poster Session	Attendees (No Government)

OFFICE OF THE DIRECTOR OF NATIONAL INTELLIGENCE

MICrONS Program Overview

LEADING INTELLIGENCE INTEGRATION

R. Jacob Vogelstein, Program Manager
IARPA Office of Safe and Secure Operations

Presentation Outline

- Introduction and motivation
- MICrONS goals and approach
- Program metrics
- Summary and feedback request

Motivation

- Brain exhibits a remarkable capacity for recognition and learning in physical and abstract data that far exceeds the capabilities of today’s state of the art machine learning systems
- Performance gap exists not only for high-level cognitive processes (e.g. “understanding”), but also for basic sensory information processing tasks supporting these higher-level functions

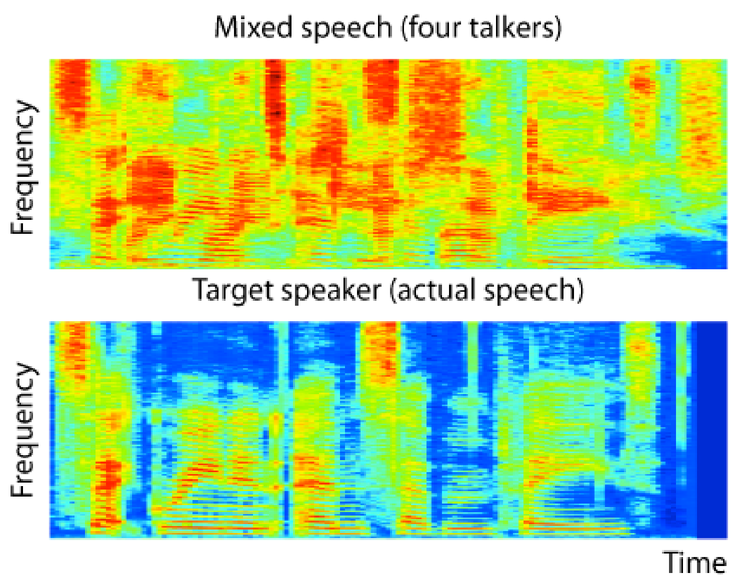

Scene Decomposition

Element Clustering

MIT SUN Database; Hershey et al., 2011; Tenenbaum et al., 2011

How Do We Do It?

- Many contemporary theories of cortical computing suggest that, for a given sensory information processing task, the brain employs algorithms composed of multiple instances of a limited set of computing *primitives*
- As defined here, primitives are:
  - Repeated structural and/or functional motifs used by one or more cortical area(s) to implement "core functions" of cortical algorithms: **representing data, transforming data, and learning from data**
  - Multiple copies reflect the need to span some sensory space and suggest some common computations performed over a feature hierarchy (or throughout the brain)
  - Constructed from $O(10^2)$–$O(10^4)$ neurons
  - Primarily local in their sites of action, i.e. largely contained within a 50–1,000 µm diameter column
- These primitives (or larger modules composed of primitives) are thought to be arranged in a multi-stage/hierarchical processing architecture with extensive feed-forward, feedback, and lateral connections within and between elements

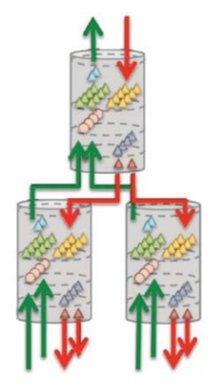

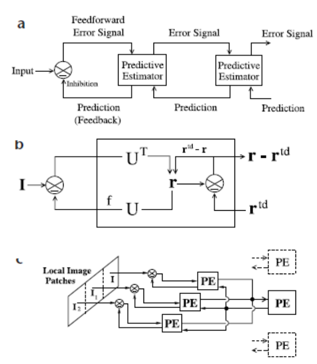

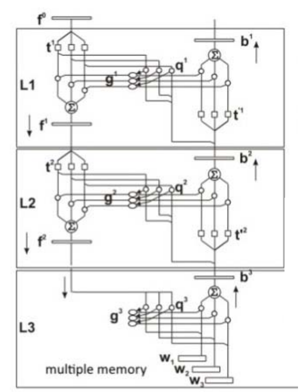

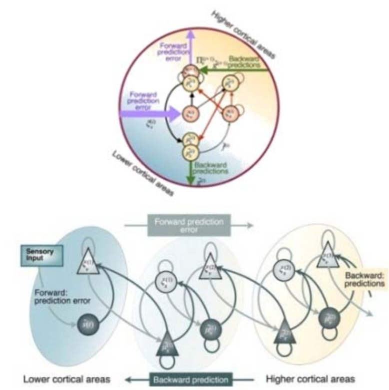

neurdon.com for Numenta, Inc.; Rao et al., 1999; Arathorn; Park and Friston, 2013

Are We There Yet?

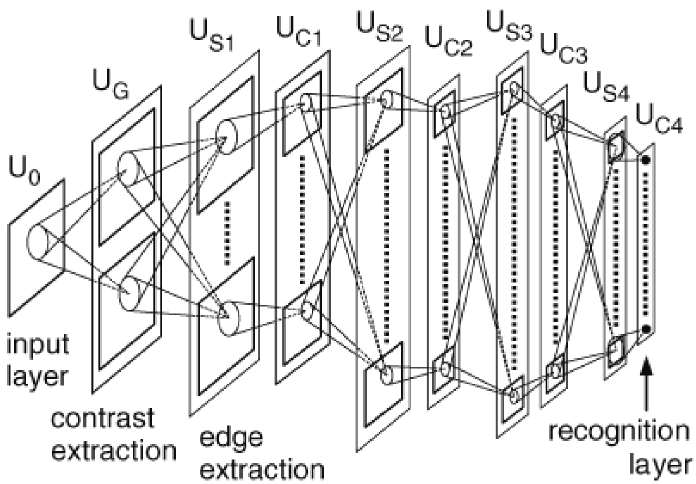

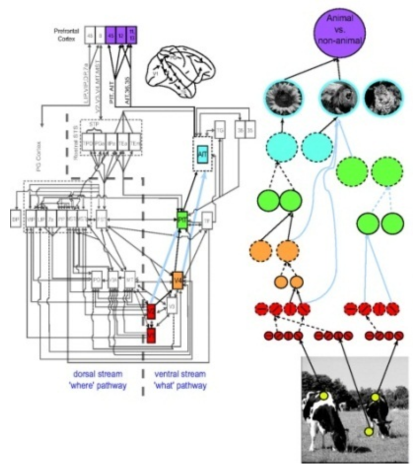

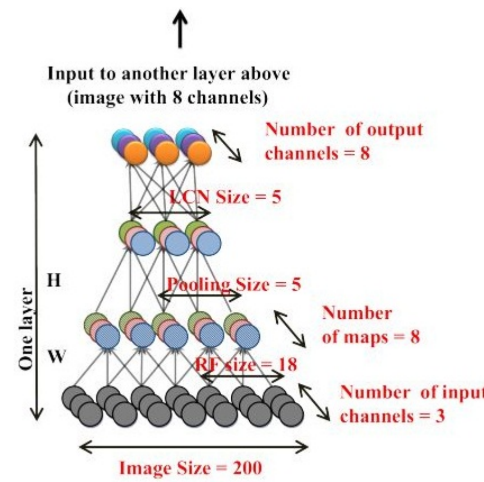

	aeroplane	bicycle	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	motorbike	person	plant	sheep	sofa	train	tv	mAP
DPM [12]	46.3	49.5	4.8	6.4	22.6	53.5	38.7	24.8	14.2	10.5	10.9	12.9	36.4	38.7	42.6	3.6	26.9	22.7	34.2	31.2	26.6
[9] + 33 context	46.2	40.5	10.5	12.4	16.5	53.2	36.2	28.2	17.2	8.3	14.5	18.8	27.9	35.4	41.9	9.3	28.7	19.9	38.6	28.6	26.7
[9] + GT 33 context	48.8	43.1	11.9	**14.3**	25.1	53.2	43.1	26.0	14.5	13.4	9.2	15.0	30.0	36.9	36.7	**10.9**	31.6	22.4	41.5	**35.5**	28.2
DPM rescoring [12] + 20 context	44.3	**51.3**	7.1	8.0	21.8	56.0	41.2	18.4	13.8	11.7	10.4	13.5	38.3	42.7	44.6	3.7	27.0	24.3	38.0	32.2	27.4
DPM rescoring [12] + 33 context	46.4	50.8	7.5	8.2	21.2	55.3	41.6	20.0	**14.7**	11.8	11.6	13.9	37.9	40.2	45.1	4.2	24.1	**27.6**	40.8	33.9	27.8
Ours + 20 context	46.9	50.1	9.2	9.5	**30.1**	**57.2**	44.1	30.7	12.7	**15.1**	12.9	14.2	35.6	**44.8**	44.0	4.9	30.6	20.1	42.2	34.8	29.5
Ours + 33 context	**49.8**	48.8	**12.0**	10.8	29.1	55.2	**45.6**	**32.0**	14.2	12.6	**13.7**	**16.6**	**39.8**	44.2	**45.1**	8.2	**35.3**	26.0	**42.3**	34.3	**30.8**

Fukushima, 1980; Serre et al., 2007; Le et al., 2012; Mottaghi et al., 2014

Why Not?

- Most of what is really known about the brain is about its microscale ($\leq O(10^0)$ neurons) and macroscale ($\geq O(10^5)$ neurons) operation and organization
- Relatively little is known about the detailed structure and function of mesoscale ($O(10^2)$–$O(10^4)$ neurons) circuits that embody the cortical computing primitives
- Implications:
  - Current approximations of neural algorithms are insufficient
  - Achieving brain-like performance requires knowledge and use of
    - Specific data **representations**, data **transformations**, and rules for **learning** employed by the brain and implemented by the cortical computing primitives
    - Top-down and lateral **feedback** connections between/within primitives

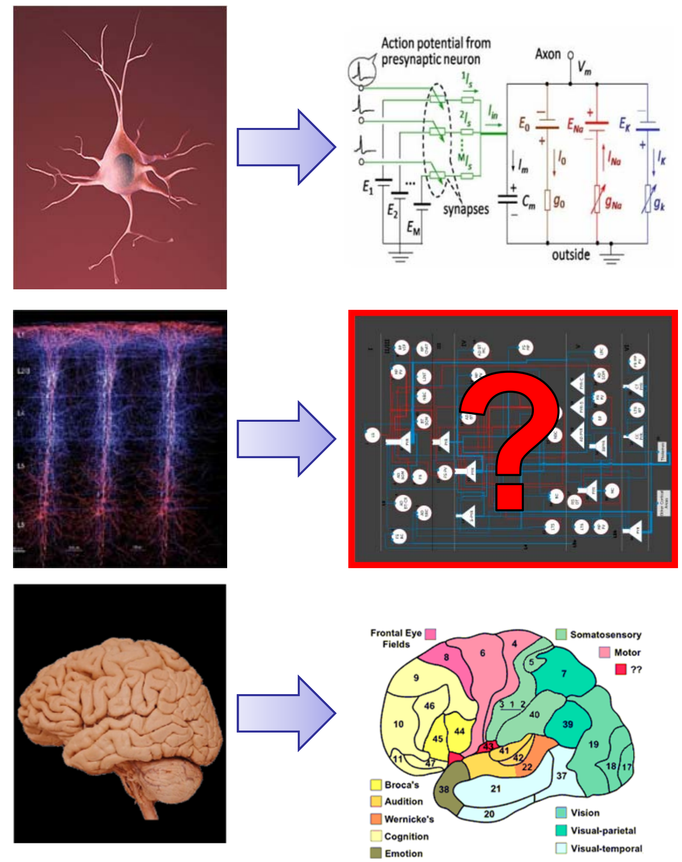

Haven't We Tried This Before?

- There have been a number of projects/programs in neurally-inspired computing that model the brain at varying scale and fidelity
- None have had access to the detailed structure and function of the mesoscale computing circuits (primitives) that implement the core functions of cortical algorithms

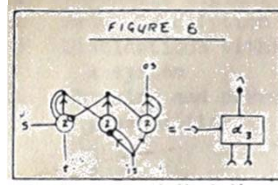

First Draft of a Report on the EDVAC

by

John von Neumann

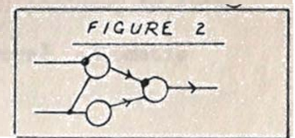

4.2 It is worth mentioning, that the neurons of the higher animals are definitely elements in the above sense. They have all-or-none character, that is two states: Quiescent and excited. They fulfill the requirements of 4.1 with an interesting variant: An excited neuron emits the standard stimulus along many lines (axons). Such a line can, however, be connected in two different ways to the next neuron: First: In an excitatory synapsis, so that the stimulus causes the excitation of that neuron. Second: In an inhibitory synapsis, so that the stimulus absolutely prevents the excitation of that neuron by any stimulus on any other (excitatory) synapsis. The neuron also has a definite reaction time, between the reception of a stimulus and the emission of the stimuli caused by it, the synaptic delay.

Following W. Pitts and W. S. MacCulloch ("A logical calculus of the ideas immanent in nervous activity", Bull. Math. Biophysics, Vol. 5 (1943), pp 115-133) we ignore the more complicated aspects

von Neumann, 1945 via Dyson, 2012

What's Changed?

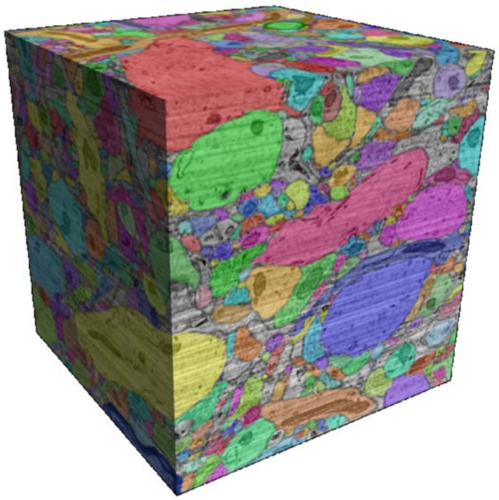

Serial Electron Microscopy

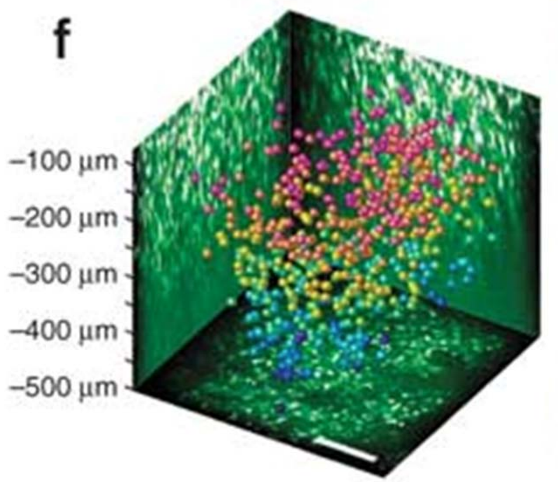

Volumetric Calcium Imaging

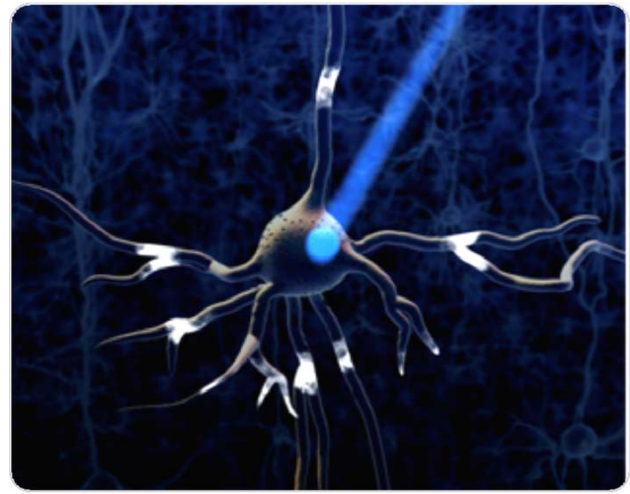

Optogenetics

...and more! Cre lines, CLARITY, array tomography, etc., etc.,

For the first time in history the necessary tools, techniques, and technologies exist to reveal *and exploit* the detailed structure, function, and networks of the cortical computing primitives

Berger; Katona et al., 2011; Boyden

Presentation Outline

- Introduction and motivation
- MICrONS goals and approach
- Program metrics
- Summary and feedback request

Program Goal

Create a new generation of machine learning algorithms that achieve human-like performance characteristics by using the same “core functions” (for representing, transforming, and learning from data) as those employed by the brain and effected by the cortical computing primitives

Overview of Approach

- MICrONS is expected to solicit proposals from vertically-integrated, multidisciplinary teams comprising
  - Theoretical and computational neuroscientists
  - Experimental neuroscientists
  - Computer scientists
  - Mathematicians and statisticians
  - Data scientists (machine learning specialists)
- Each team is expected to:
  1. Posit one or more algorithmic framework(s) to describe sensory information processing in a given region(s) of the brain and describe how the framework(s) represents data, transforms data, and learns from data (i.e. identify candidate primitives)
  2. Collect and analyze <u>detailed</u> structural and functional data on mesoscale cortical circuits (~1 mm^3 at nm/dsec resolution) in **mammalian** brain(s) that can be used in computational neural models to inform, refine, (dis)prove, and/or disambiguate various aspects of the primitives and frameworks
  3. Develop computational neural models consistent with this new data and the existing literature to elucidate the architecture and operation of the cortical computing primitives
  4. Instantiate the frameworks in the form of machine learning algorithms whose core functions (representations, transformations, and learning rules) are supported by and consistent with computational neural models of cortical computing primitives

Approach in Detail (#1)

Algorithmic framework for cortical information processing

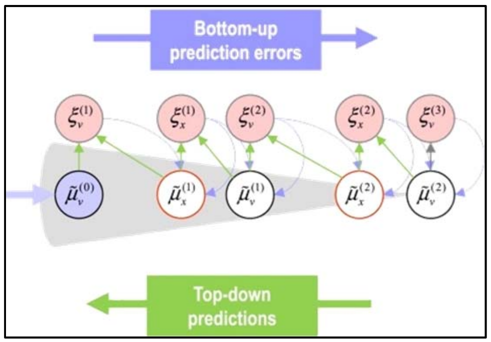

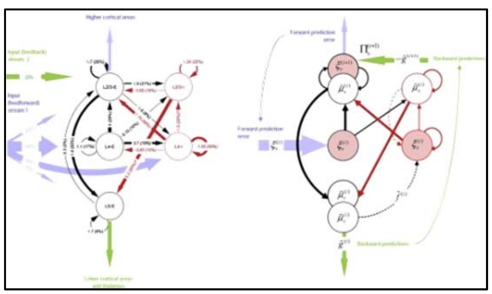

Neurally-plausible implementation of core functions

- For each framework, the team should specify:
  - A high-level “objective function”
  - How the framework **represents** data (e.g. deterministic vs. probabilistic, sparse vs. dense coding)
  - How the framework **transforms** data during processing, i.e. dynamics of computation (e.g. filtering, sampling, etc.)
  - The rules that govern supervised and unsupervised **learning** in the framework (e.g. locally-driven learning, propagated learning signals, etc.)
  - How top-down and lateral **feedback** are used during perception and learning (e.g. gain modulation, feature selectivity, normalization, etc.)
  - Initial candidates of plausible neural mechanisms (i.e. primitives) realizing these core functions
- Constraints:
  - The framework should be compatible with existing neuroscience data and biophysical processes
  - The framework should have characteristics that suggest it will solve challenging machine learning tasks differently and, ultimately, better than today’s state of the art algorithms
  - The framework should be difficult to fully realize without additional data about the structure and function of cortical microcircuits (i.e. we’re not looking to simply confirm theory, but rather to *inform* theory through targeted experiments and data collection)

Bastos et al., 2012

Approach in Detail (#2)

Algorithmic framework for cortical information processing

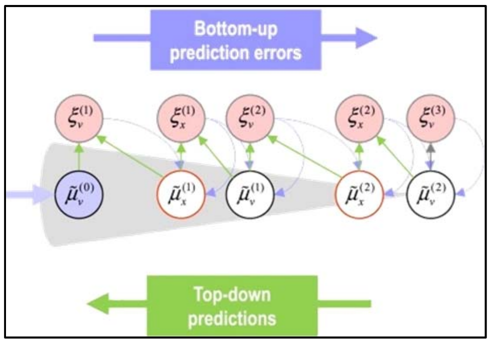

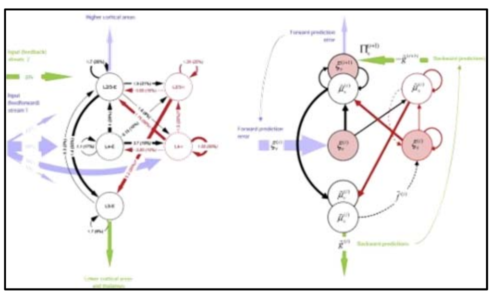

Neurally-plausible implementation of core functions

BAA Response

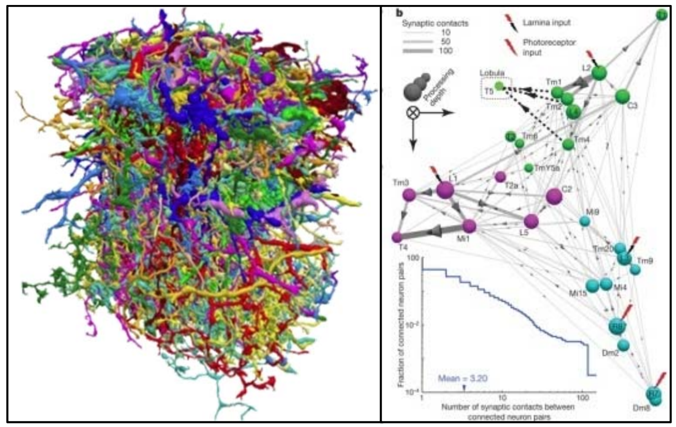

Directed data collection and analysis

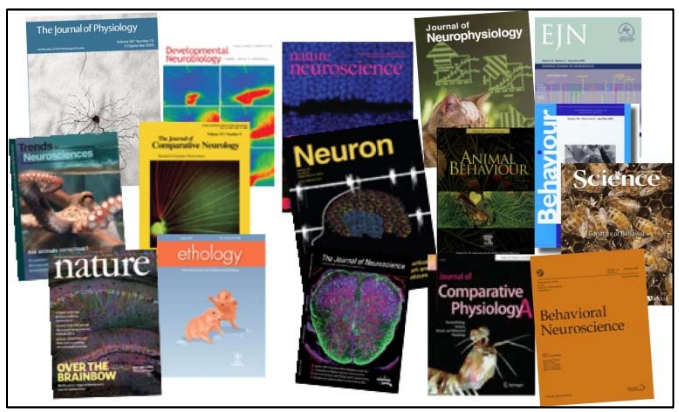

- What data about the structure and function of mesoscale cortical circuits can be collected to elucidate the framework and its associated core functions and cortical computing primitives?
  - Which cortical area(s)?
  - Which animal model(s)?
  - Which behavior(s)?
  - Which experimental techniques?
- What other data (beyond ~1 mm^3) can be collected to inform the implementation of the framework?
- Remember: interesting ≠ useful

Bastos et al., 2012; Chklovskii; Takemura et al., 2013

Approach in Detail (#3)

Algorithmic framework for cortical information processing

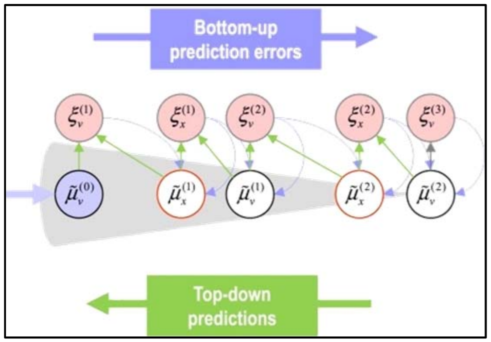

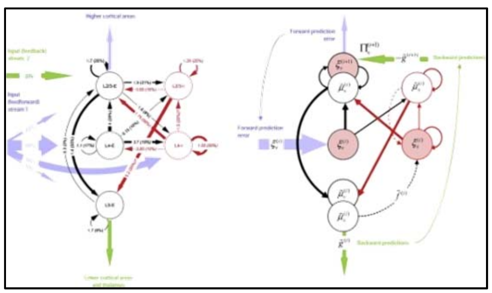

Neurally-plausible implementation of core functions

BAA Response

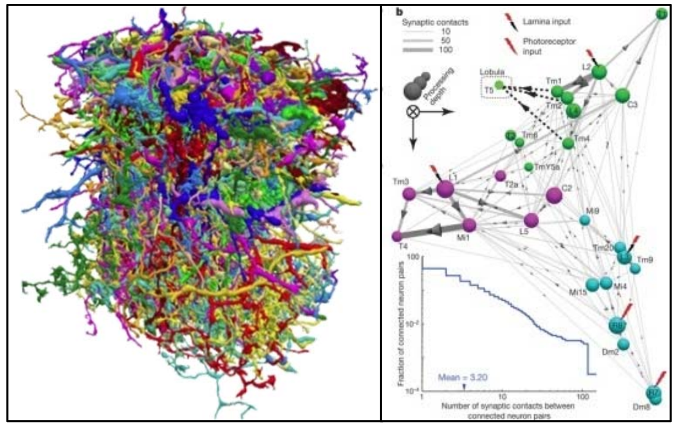

Directed data collection and analysis

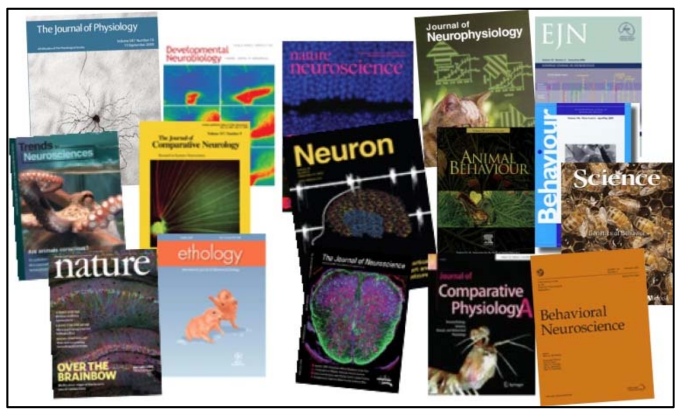

Computational neural models

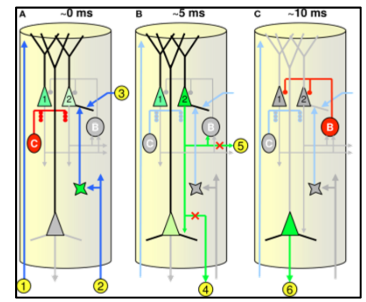

- Models should:
  - Help to ascertain the functions of the cortical computing primitives
  - Constrain the design of machine learning algorithms: algorithms' core functions should be supported by observed and modeled brain data

Bastos et al., 2012; Chklovskii; Takemura et al., 2013; Rinkus, 2010

Approach in Detail (#4)

Algorithmic framework for cortical information processing

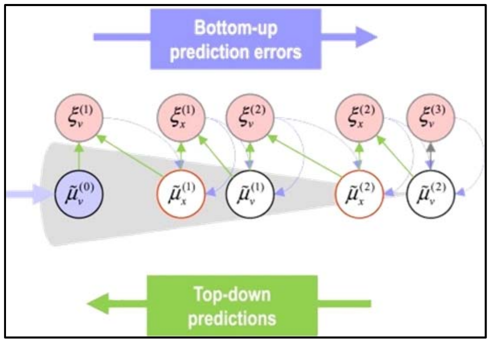

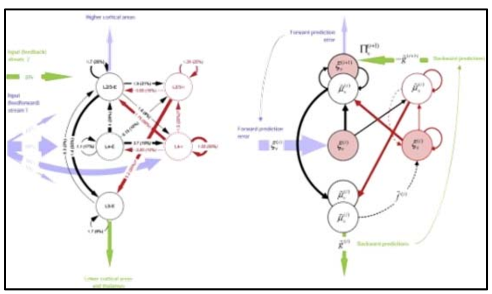

Neurally-plausible implementation of core functions

BAA Response

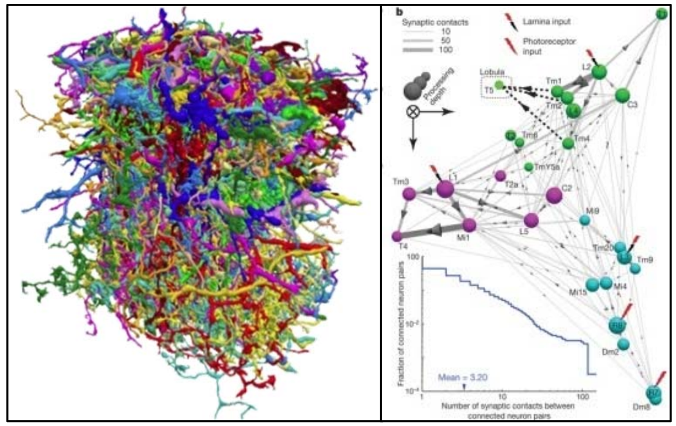

Directed data collection and analysis

- Machine learning algorithms should:
  - Use representations, transformations, and learning rules (i.e. core functions) that are consistent with the cortical computing primitives
  - Incorporate top-down and lateral feedback during perception and learning
  - Exhibit “human-like” performance on sensory information processing tasks, such as sensory “scene parsing”
  - Generalize to abstract, non-sensory data

Computational neural models

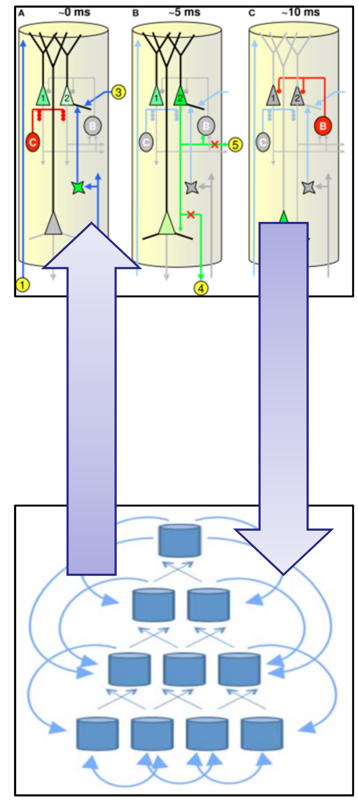

Machine learning algorithms

Bastos et al., 2012; Chklovskii; Takemura et al., 2013; Rinkus, 2010

Program Timeline & Milestones

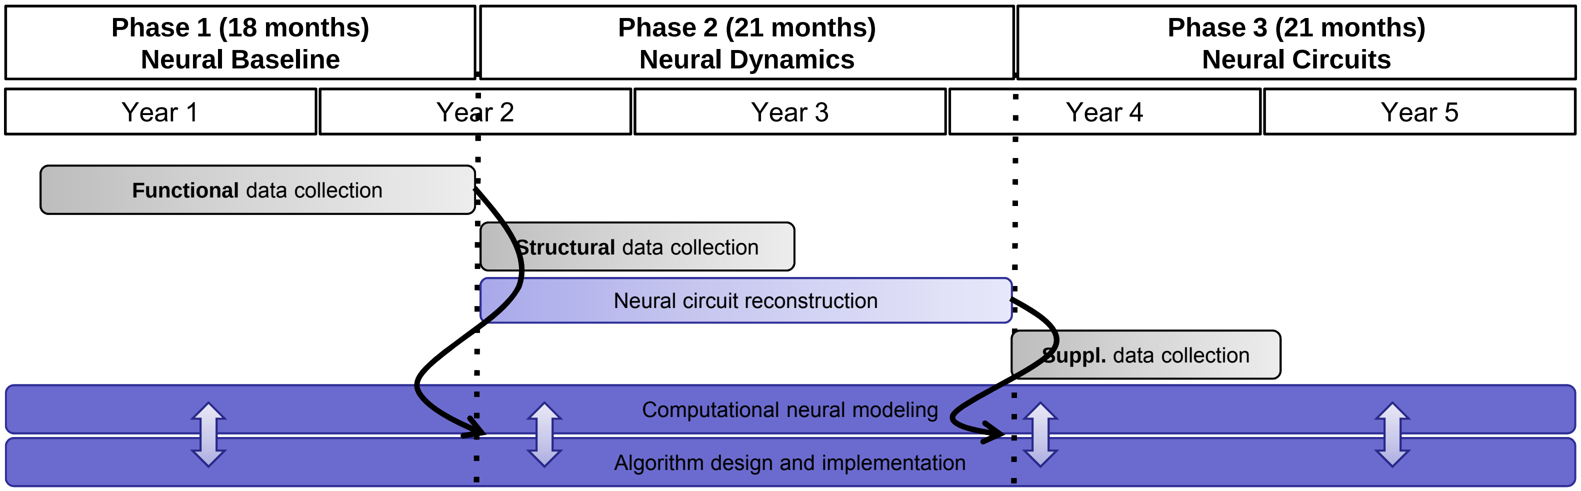

- Baseline algorithm
  - Proof-of-concept implementation of algorithmic framework with biologically-plausible neural architecture, components, and learning rules
  - Core functions supported by computational neural models based on existing literature
  - Validated by ≥75% performance in parsing simple synthetic scenes

- Updated algorithm
  - Functional data constrains updated models and informs updated algorithm design
  - Algorithms exploit observed neural representations and dynamics to enhance performance
  - Validated by ≥75% performance in parsing more realistic synthetic scenes

- Final algorithm
  - Exploits observed function and structure of cortical computing primitives to maximize biological fidelity of models and algorithms and (presumably) maximize performance
  - Validated by ≥75% performance in parsing real-world (recorded) scenes
  - Demonstrate generalizability of core functions to abstract, non-sensory data

Alternative Timeline & Milestones

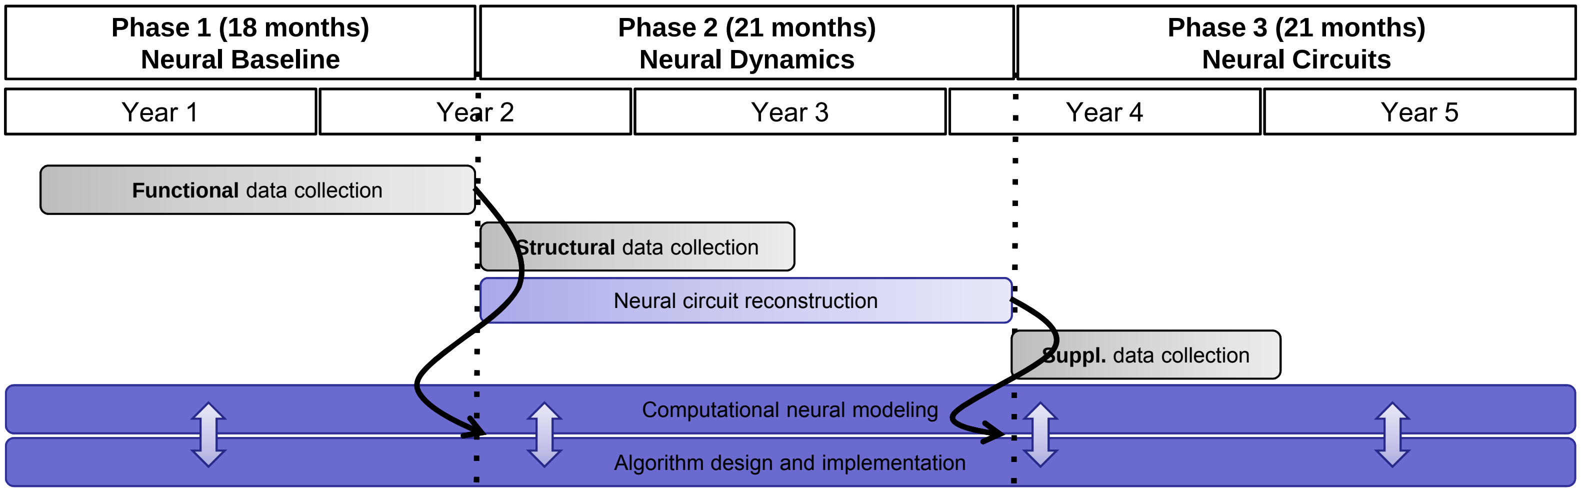

- Baseline algorithm
  - Proof-of-concept implementation of algorithmic framework with biologically-plausible neural architecture, components, and learning rules
  - Core functions supported by computational neural models based on existing literature
  - Verify piecewise operation of core functions and correspondence with associated computational neural models
- Updated algorithm
  - Functional data constrains updated models and informs updated algorithm design
  - Algorithms exploit observed neural representations and dynamics to enhance performance
  - Validated by ≥75% performance in parsing simple synthetic scenes
- Final algorithm
  - Exploits observed function and structure of cortical computing primitives to maximize biological fidelity of models and algorithms and (presumably) maximize performance
  - Validated by ≥75% performance in parsing more complex synthetic scenes
  - Demonstrate generalizability of core functions to abstract, non-sensory data

Presentation Outline

- Introduction and motivation
- MICrONS goals and approach
- Program metrics
- Summary and feedback request

Areas of Innovation and Assessment

- MICrONS will seek innovations in the following areas:
  - Neural data acquisition: collecting "large scale" co-registered structural and functional datasets to interrogate mesoscale cortical circuits (at nanometer and sub-second resolution)
  - Neural circuit reconstruction: converting raw data into annotated schematics describing the underlying neural circuits and identifying repeated structural/functional motifs that comprise the cortical computing primitives
  - Computational neural modeling: creating computational models of the observed structural and functional data to explicate the cortical computing primitives and to constrain the design of novel machine learning algorithms
  - Machine learning algorithms: designing and implementing novel machine learning algorithms that use the same "core functions" as those employed by the brain and effected by the cortical computing primitives
  - *Theoretical neuroscience: understanding the principles of sensory information processing and cortical computing*

Areas of Innovation and Assessment

- MICrONS will seek innovations in the following areas:
  - Neural data acquisition: collecting "large scale" co-registered structural and functional datasets to interrogate mesoscale cortical circuits (at nanometer and sub-second resolution)
  - Neural circuit reconstruction: converting raw data into annotated schematics describing the underlying neural circuits and identifying repeated structural/functional motifs that comprise the cortical computing primitives
  - Computational neural modeling: creating computational models of the observed structural and functional data to explicate the cortical computing primitives and to constrain the design of novel machine learning algorithms
  - Machine learning algorithms: designing and implementing novel machine learning algorithms that use the same "core functions" as those employed by the brain and effected by the cortical computing primitives
  - *Theoretical neuroscience: understanding the principles of sensory information processing and cortical computing*
- Metrics are designed to assess whether intermediate and final work products in each of these areas justify continued investment

Machine Learning Algorithm Metrics

- Machine learning algorithms will be evaluated on two criteria:
  1. Neural fidelity
     - Definition: Correspondence between the "core functions" of the algorithm and the associated neural models
     - Evaluation: Defined collaboratively between performers and T&E team – similar to approach for computational neural models (see below)
     - Target: "Pass"
  2. Performance on two "scene parsing" tasks
     A. Scene decomposition
        - Definition: Isolation (aka segmentation) of individual scene elements (visual objects, sound sources, etc.)
        - Evaluation: "Pixel"-level segmentation accuracy per element
        - Target: 75% accuracy (in all phases)

     B. Scene element clustering
        - Definition: Grouping of similar elements (in sensory feature space) within and across scenes
        - Evaluation: Cluster quality relative to human manual clustering, as determined by Adjusted Rand Index or similar measure
        - Target: 75% accuracy (in all phases)
     - Challenge problems (aka scenes) increase in complexity over the three program phases
       - Phase 1: "Simple" synthetic scenes
       - Phase 2: "Complex" synthetic scenes
       - Phase 3: Real (recorded) scenes
- In all phases, challenge problems will be modality-specific (i.e. algorithms derived from visual cortex will be tested on visual scenes, algorithms derived from olfactory cortex will be tested on olfactory scenes, etc.)

Easy Visual Scene Parsing Illustration (Phase 1)

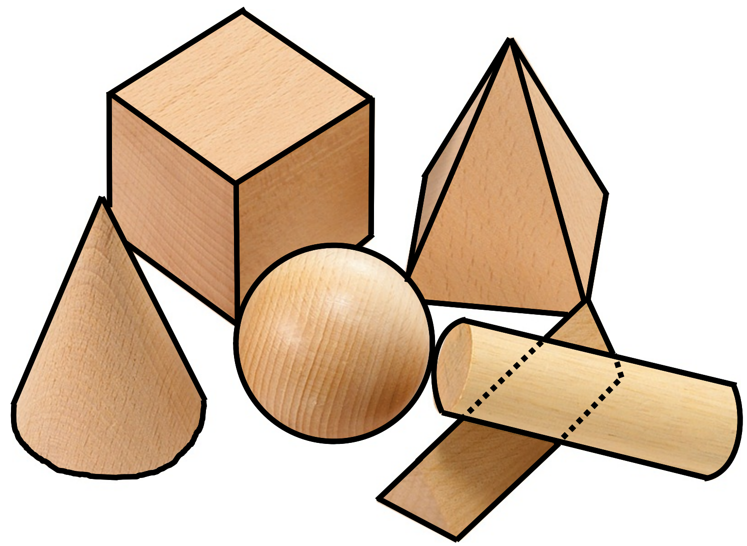

Emergent cluster across images (cylinders)

Hard Visual Scene Parsing Illustration (Phase 3)

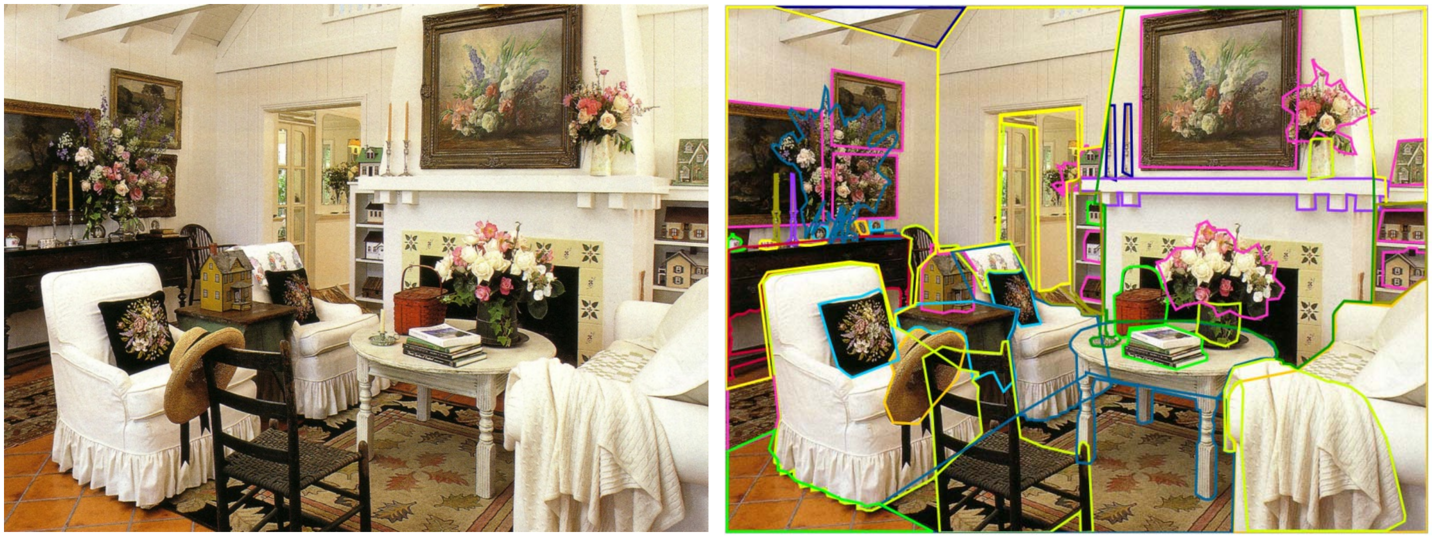

Emergent cluster within image (houses)

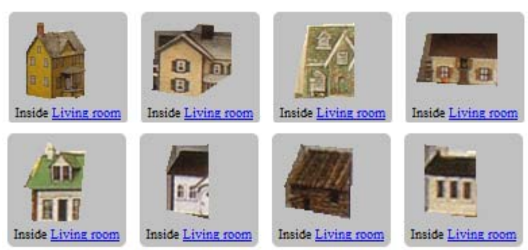

MIT SUN Database

Auditory Scene Parsing Illustration

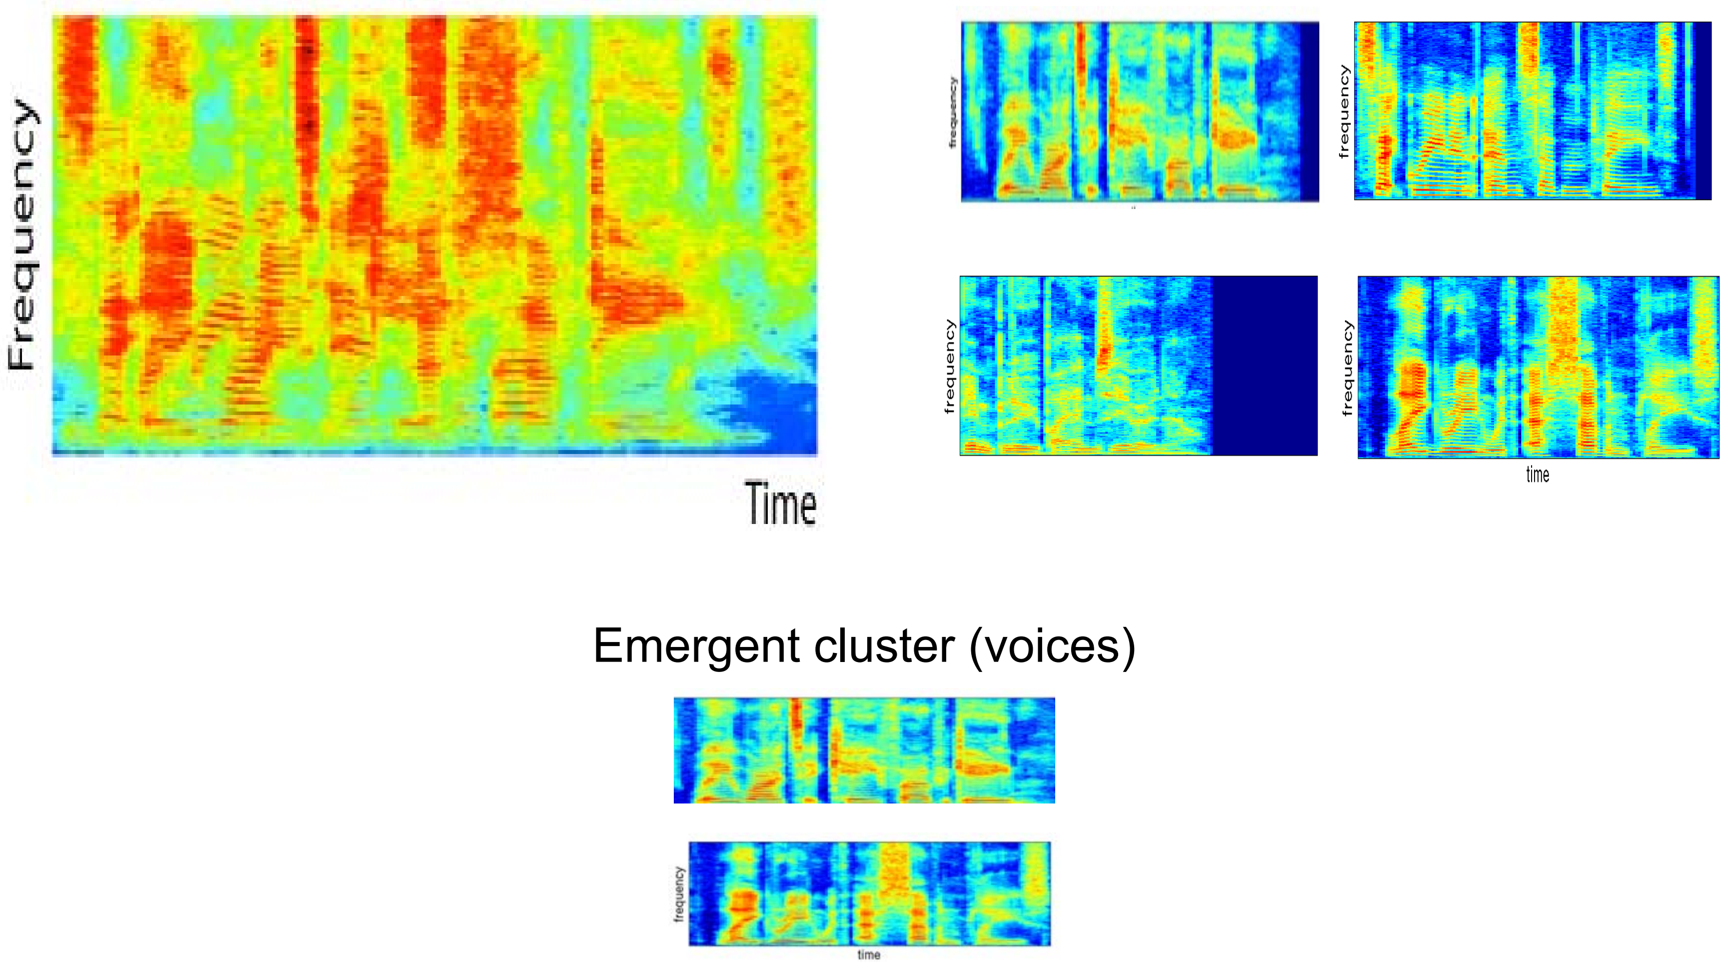

Hershey et al., 2011; Rennie et al., 2010

Olfactory Scene Parsing Illustration

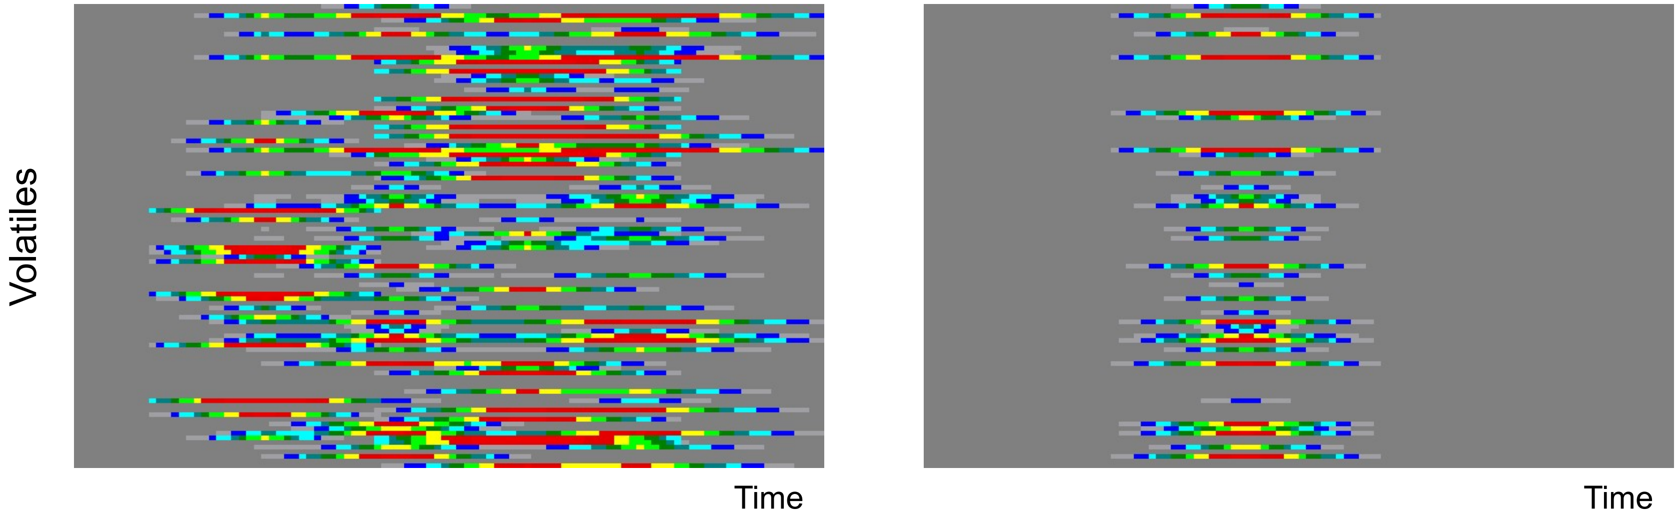

Emergent cluster (citrus)

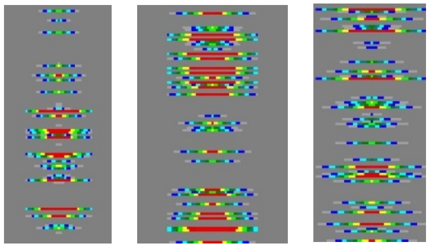

Rhodes

Why Scene Parsing?

- Scene parsing has many applications in the intelligence community, but **this program is not about scene parsing** per se
- Rationale for scene parsing
  - MICrONS is seeking insights into the core functions of cortical algorithms (the way the brain represents data, transforms data, and learns from data) to improve machine learning
  - To exercise the discovered functions, I believe they must be embedded in an algorithm that actually “does something” in machine learning
  - If the task the algorithm performs is too easy, it may not demonstrate the limits (capabilities) of the algorithm and the value of using neurally-derived functions versus conventional approaches
  - If the task the algorithm performs is too complex or too specific, it may require too much engineering or training that is tangential to exercising the core functions
  - Scene parsing is perhaps the “simplest hard problem” in sensory information processing that is accessible to currently-available experimental techniques and to currently-available conventional machine learning algorithms
    - It is easy for the brain and likely engages core functions at multiple stages of information processing
    - It is hard for conventional machine learning algorithms
  - The two scene parsing tasks defined here have a number of nice computational attributes
    - They do not require storing extensive vocabularies of scene elements (nothing has to be labeled)
    - They probably collectively involve some supervised and mostly unsupervised learning
    - They likely require feedback/recurrence (to resolve sensory ambiguity in complex scenes)
    - They highlight solutions to two fundamental problems that transcend the data domain
      - Resolving ambiguity
      - Identifying salient features
  - Scene parsing also lends itself to a well-defined and replicable third-party evaluation framework
- And/but: If you have a better (or equally good) idea, I want to hear it!

Computational Neural Models Metrics

- Computational neural models will be evaluated on three criteria:
  1. Structural fidelity
  2. Functional fidelity
  3. Model credibility
- How do we assign a number to these metrics?
  – We don't; these are Pass/Fail assessments
  – A MICrONS T&E team will be established to:
    - Work with performers to collaboratively define a specific methodology for assessing the structural and functional fidelity
      – Structural fidelity may include aspects like the number of types of cells modeled and the connectivity or morphology of cells modeled
      – Functional fidelity may include correspondence between modeled and observed outputs, given similar (simulated) inputs
    - Convene a panel of experts to assess models' overall credibility
      – Credibility may include consideration of factors such as:
        » Robustness to parameter variation
        » Economy of explanation
        » Reliance upon uncommon or unlikely neural mechanisms or parameter values

Data Collection Metrics Overview

- Minimum specifications will be established on three aspects of the structural and functional data to be collected on MICrONS:
  - Spatial extent
    - Structural data: ≥ 1 mm^3 overall, spanning full cortical thickness
    - Functional data: ≥ Cortical surface area of structural images x 500 um^3
  - Spatial resolution
    - Structural data: ≤ 5 x 5 x 30 nm^3 (smaller is better)
    - Functional data: ≤ 5 x 5 x 12.5 um^3 (smaller is better)
  - Temporal resolution
    - Functional data: ≥ 4 Hz (per neuron, not necessarily simultaneous)
- Performers will also propose additional metrics (and target values) appropriate for their particular technical approach, such as:
  - Image contrast
  - Missing voxels
  - Neuron coverage (% neurons recorded)

Circuit Reconstruction Metrics Overview

- Circuit reconstruction will be evaluated on three criteria:
  1. Overall accuracy of connectivity
     - Definition: Agreement between the reconstructed (R) and ground-truth (G) *weighted* adjacency matrices
     - Evaluation: MICrONS T&E will (semi-)manually reconstruct some subset of structural images to create G and compute row-wise average agreement between R and G as $\aleph = (R \cdot G) \div |R||G|$
     - Target: $\aleph$ @ $\{10\%ile, 50\%ile, 90\%ile\} \geq \{0.05, 0.70, 0.95\}$
  2. Connection specificity
     - Definition: Number of non-connections between neurons correctly specified divided by the total number of non-connections in the reconstruction
     - Evaluation: MICrONS T&E will threshold and binarize R & G matrices and compute average specificity over sampled portions of the matrices
     - Target: ≥ 0.90
  3. Connection sensitivity
     - Definition: Number of connections between neurons correctly identified divided by the total number of connections detected
     - Evaluation: MICrONS T&E will threshold and binarize R & G matrices and compute average sensitivity over sampled portions of the matrices
     - Target: ≥ 0.90
- Performers will propose additional metrics on other aspects of circuit reconstruction such as functional and structural registration, circuit motifs, morphology of constituent cells, etc.

Technical Milestones

Milestone 1 **Phase 1 / Month 3**	• All IACUC protocols approved.
Milestone 2 **Phase 1 / Month 18**	• Data: Complete functional dataset delivered; all functional data acquisition target values achieved. • Models: Baseline model description and source code delivered; passing score on model credibility. • Algorithms: Executable, performance report, and documented source code delivered; all Phase 1 scene parsing target values achieved.
Milestone 3 **Phase 2 / Month 27**	• Data: Complete structural dataset delivered; all structural data acquisition target values achieved.
Milestone 4 **Phase 2 / Month 39**	• Circuit reconstruction: Weighted and binarized adjacency matrices delivered; all circuit reconstruction target values achieved. • Models: Intermediate model description and source code delivered; high functional fidelity and passing model credibility scores achieved. • Algorithms: Executable, performance report, and documented source code delivered; all Phase 2 scene parsing target values achieved; proposed generalization demonstration documented.
Milestone 5 **Phase 3 / Month 48**	• Data: Updated structural and functional data (as necessary) delivered.
Milestone 6 **Phase 3 / Month 60**	• Circuit reconstruction: Updated adjacency matrices (as necessary) delivered. • Models: Final model description and source code delivered; high functional fidelity, high structural fidelity, and passing model credibility scores achieved. • Algorithms: Executable, performance report, and documented source code delivered; all Phase 3 scene parsing target values achieved; algorithm generalization demonstrated.

Presentation Outline

- Introduction and motivation
- MICrONS goals and approach
- Program metrics
- Summary and feedback request

Summary

- Goal:
  - Create a new generation of machine learning algorithms that use the same representations, transformations, and learning rules as those employed by the brain and effected by the cortical computing primitives
  - Validate the algorithms through performance of machine learning tasks on sensory data and illustration of generalization of core functions to abstract, non-sensory data
- Approach
  - Form vertically-integrated, multidisciplinary teams
  - Posit algorithmic frameworks that describe sensory information processing in the brain
  - Collect and analyze structural and functional data on mesoscale cortical circuits in mammalian brain(s)
  - Develop computational neural models consistent with this new data and the existing literature to elucidate the architecture and operation of the cortical computing primitives
  - Instantiate the frameworks in the form of machine learning algorithms whose core functions are supported by and consistent with computational neural models of cortical computing primitives

Feedback Request

- We want to increase your chances of success in achieving MICrONS goals
- We are specifically seeking your input on the following questions regarding the machine learning task and associated metrics
  - Are algorithms for “scene decomposition” and “scene element clustering”:
    1. Likely to exercise the core functions of cortical algorithms?
    2. Likely to be informed by interrogating sensory cortex at the resolution and scale specified previously (i.e. 1 mm^3 at nanometer and sub-second resolution)?
    3. Likely to depend primarily on other parts of the brain, such that it would be more appropriate for MICrONS to focus on a subcomponent or subtask within scene parsing?
    4. Too hard (in Phase 1, 2 and/or 3) to be useful gauges of progress?
    5. Too easy to motivate interest or to demonstrate unique capabilities of the novel algorithms?
    6. Dependent on too much engineering tangential to the core functions of interest?
  - Are the “scene decomposition” and “scene element clustering” tasks:
    1. Over- or under-constrained?
    2. Too narrow to exploit neuroscience findings?
  - Are there parts of the brain you’d want to study that are compatible with the rest of the program concept, but unlikely to inform algorithms for scene parsing?
    - Which tasks would be more appropriate for these brain areas?
    - What if the task was restricted to auditory or visual (and not, e.g., olfactory) scene parsing?
  - What alternative tasks would you suggest that are better suited to your approach?
- If you have answers to any of these questions, please submit an index card!

Proposers' Day Agenda

8:00AM – 8:30AM	Registration	
8:30AM – 8:35AM	Welcome and opening remarks	R. Jacob Vogelstein Program Manager
8:35AM – 9:00AM	IARPA Overview	Peter Highnam Director, IARPA
9:00AM – 10:00AM	MICrONS Program Overview	R. Jacob Vogelstein Program Manager
10:00AM – 10:30AM	**Break**	
10:30AM – 11:00AM	Doing Business with IARPA	Tarek Abboushi IARPA Acquisitions
11:00AM – 12:00PM	MICrONS Program Feedback and Q&A	R. Jacob Vogelstein Program Manager
12:00PM – 1:30PM	Lunch / Poster Session	
1:30PM – 4:00PM	Proposers' Capabilities Briefings	Attendees (No Government)
4:00PM – 5:00PM	Poster Session	Attendees (No Government)

OFFICE OF THE DIRECTOR OF NATIONAL INTELLIGENCE

Doing Business with IARPA

LEADING INTELLIGENCE INTEGRATION

July 17, 2014

Doing Business with IARPA Recurring Questions

- Questions and answers (http://www.iarpa.gov/index.php/faqs)
- Eligibility information
- Intellectual property
- Pre-publication review
- Preparing the proposal (Broad Agency Announcement (BAA) Section 4)
  - Electronic proposal delivery (https://iarpa-ideas.gov)
  - Organizational Conflicts of Interest (OCOI)
  - http://www.iarpa.gov/index.php/working-with-iarpa/iarpas-approach-to-oci
- Streamlining the Award Process
  - Accounting system
  - Key personnel
- IARPA funds "applied research"
- RECOMMENDATION: Read the entire BAA

Responding to Q&As

- Please read entire BAA before submitting questions
- Pay attention to Section 4 (Application & Submission Info)
- Read Frequently Asked Questions on the IARPA website @ http://www.iarpa.gov/index.php/faqs
- Send your questions as soon as possible
  - MICRONS BAA: dni-iarpa-baa-14-06@iarpa.gov
  - Write questions as clearly as possible
  - Do NOT include proprietary information

Eligible Applicants

- Collaborative efforts/teaming strongly encouraged
  - Content, communications, networking, and team formation are the responsibility of Proposers
- Foreign organizations and/or individuals may participate
  - Must comply with Non-Disclosure Agreements, Security Regulations, Export Control Laws, etc., as appropriate, as identified in the BAA

Ineligible Organizations

- Any organizations that have a special relationship with the Government, including access to privileged and/or proprietary information, or access to Government equipment or real property, to include:
  - Other Government Agencies
  - Federally Funded Research and Development Centers (FFRDCs)
  - University Affiliated Research Centers (UARCs)
- Are **NOT** eligible to submit proposals under this BAA or participate as team members under proposals submitted by eligible entities

Intellectual Property (IP)

- Unless otherwise requested, Government rights for data first produced under IARPA contracts will be UNLIMITED
- At a minimum, IARPA requires Government Purpose Rights (GPR) for data developed with mixed funding
- Exceptions to GPR: State in the proposal any restrictions on deliverables relating to existing materials (data, software, tools, etc.)
- If selected for negotiations, you must provide the terms relating to any restricted data or software, to the Contracting Officer

Pre-Publication Review

- For funded applied research efforts, IARPA encourages publication for peer review of UNCLASSIFIED research
- Prior to public release of any work submitted for publication, the Performer will:
  - Provide copies to the IARPA PM and Contracting Officer Representative (COR/COTR)
  - Ensure shared understanding of applied research implications between IARPA and Performers

Preparing the Proposal

- Note restrictions in BAA Section 4 on proposal submissions
  – Interested Offerors must register electronically in accordance with instructions on https://iarpa-ideas.gov
  – Interested Offerors are strongly encouraged to register in IDEAS at least 1 week prior to proposal “Due Date”
  – Offerors must ensure the version submitted to IDEAS is the “Final Version”
  – Classified proposals: Contact IARPA Chief of Security using the contact info provided in the BAA
- BAA format is established to answer most questions
- Check FBO for amendments and IARPA website for Q&As
- BAA Section 5 – Read Evaluation Criteria carefully (e.g. “The technical approach is credible, and includes a clear assessment of primary risks and a means to address them”)

Preparing the Proposal (BAA Sect 4)

- Read IARPA's Organizational Conflict of Interest (OCI) policy: http://www.iarpa.gov/index.php/working-with-iarpa/iarpas-approach-to-oci
- See also eligibility restrictions on use of Federally Funded Research and Development Centers, University Affiliated Research Centers, and other similar organizations that have a special relationship with the Government
  - Focus on possible OCIs of your institution as well as the personnel on your team
  - See Section 4: It specifies the non-Government (e.g., SETA, FFRDC, UARC, etc.) support we will be using. If you have a potential or perceived conflict, request waiver as soon as possible

Organizational Conflict of Interest (OCI)

- If a prospective offeror, or any of its proposed subcontractor teammates, believes that a potential conflict of interest exists or may exist (whether organizational or otherwise), the offeror should promptly raise the issue with IARPA and submit a waiver request by e-mail to the mailbox address for this BAA at dni-iarpa-baa-14-06@iarpa.gov
- A potential conflict of interest includes but is not limited to any instance where an offeror, or any of its proposed subcontractor teammates, is providing either scientific, engineering and technical assistance (SETA) or technical consultation to IARPA. In all cases, the offeror shall identify the contract under which the SETA or consultant support is being provided
- Without a waiver from the IARPA Director, neither an offeror, nor its proposed subcontractor teammates, can simultaneously provide SETA support or technical consultation to IARPA and compete or perform as a Performer under this solicitation

Streamlining the Award Process

- Cost Proposal – we only need what we ask for in BAA
- Approved accounting system needed for Cost Reimbursable contracts
  - Must be able to accumulate costs on job-order basis
  - DCAA (or cognizant auditor) must approve system
  - See http://www.dcaa.mil, "Audit Process Overview - Information for Contractors" under the "Guidance" tab
- Statements of Work (format) may need to be revised
- Ensure that "Key Personnel" meet expectations of time devoted to the project; note the Evaluation Criteria requiring relevant experience and expertise
- Following selection, Contracting Officer may request your review of subcontractor proposals

IARPA Funding

- IARPA funds applied research for the Intelligence Community (IC)
  - IARPA cannot waive the requirements of Export Administrative Regulation (EAR) or International Traffic in Arms Regulation (ITAR)
  - Not subject to DoD funding restrictions for R&D related to overhead rates
- IARPA is <u>not</u> DoD

Disclaimer

- This is applied research for the Intelligence Community
- Content of the Final BAA will be specific to this program
  - The Final BAA is being developed
  - Following issuance, look for Amendments and Q&As
  - There will likely be changes
- The information conveyed in this brief and discussion is for planning purposes and is subject to change prior to the release of the Final BAA

Proposers' Day Agenda

8:00AM – 8:30AM	Registration	
8:30AM – 8:35AM	Welcome and opening remarks	R. Jacob Vogelstein Program Manager
8:35AM – 9:00AM	IARPA Overview	Peter Highnam Director, IARPA
9:00AM – 10:00AM	MICrONS Program Overview	R. Jacob Vogelstein Program Manager
10:00AM – 10:30AM	Break	
10:30AM – 11:00AM	Doing Business with IARPA	Tarek Abboushi IARPA Acquisitions
11:00AM – 12:00PM	**MICrONS Program Feedback and Q&A**	**R. Jacob Vogelstein** **Program Manager**
12:00PM – 1:30PM	Lunch / Poster Session	
1:30PM – 4:00PM	Proposers' Capabilities Briefings	Attendees (No Government)
4:00PM – 5:00PM	Poster Session	Attendees (No Government)

OFFICE OF THE DIRECTOR OF NATIONAL INTELLIGENCE

MICrONS Feedback and Q&A

LEADING INTELLIGENCE INTEGRATION

Preemptive FAQ (Part 1)

Q: What is an example of an "algorithmic framework"?

A: Examples include some forms of predictive coding, map seeking circuits, hierarchical Bayesian inference, free-energy minimization, hierarchical temporal memory, Adaptive Resonance Theory, reservoir networks, Helmholtz machines, nonlinear adaptive control, and MANY others. This is <u>NOT</u> an exhaustive list and inclusion in this list is <u>NOT</u> an endorsement; the examples listed here were selected only because they have short (2-3 word) titles convenient for PowerPoint =)

Q: Can one proposal include multiple algorithmic frameworks?

A: Yes, assuming the same experimental plan and data can be used to inform, compare, or select between the different frameworks.

Q: Is MICrONS expecting to find complete algorithms in 1 mm^3 of brain tissue?

A: No, it is seeking insights about the architectures and core functions (representations, transformations, and learning rules) used in cortical algorithms and it is expecting that machine learning algorithms that use the same architectures and functions (within a larger engineered system) will have desirable performance characteristics on many tasks.

Q: Are "standard" back-propagation neural networks a viable algorithmic framework?

A: Probably not, because MICrONS frameworks must have a role for top-down and lateral feedback during perception and must have biologically-plausible learning rules, among other similar reasons.

Preemptive FAQ (Part 2)

Q: What kind of training and testing data will be provided for the challenge problems in scene parsing?

A: Training data for visual scene parsing problems may be video or RBG-D images; training data for auditory scene parsing problems will likely be recordings from multiple microphones and/or multi-track recordings; training data for olfactory scene parsing problems could be measurements of volatiles from multiple locations; etc. Testing data in all modalities will be "flattened" representations of similar scenes.

Q: Are the algorithms expected to use simulated neurons?

A: We expect that most algorithms will be abstracted from their neural substrate but there is no reason why an algorithm couldn't use simulated neurons.

Q: Do the neural models have to use spiking neurons?

A: Not necessarily, but models will be evaluated on their functional and structural fidelity and their overall credibility, so aggregations of neurons and/or spikes must be justifiable and decomposable into their constituent individual components.

Preemptive FAQ (Part 3)

Q: Can one propose to study something other than sensory neocortex?

A: Yes, but our working plan is to validate algorithms on sensory scene parsing tasks, so a case must be made for either (1) how algorithms derived from other parts of neocortex or allocortex will generalize to these tasks, or (2) an alternative evaluation framework.

Q: Can one propose to study mouse brains? Rat brains? Cat brains? Monkey brains? Human brains?

A: Yes, any mammal is acceptable if it is appropriate for your particular technical approach, but there must be a compelling case for using non-human primates if they are to be sacrificed in service of this program.

Q: Can one propose to study a part of a non-mammalian vertebrate or invertebrate?

A: Probably not, unless it is a justifiable addition to a research plan that also incorporates mammalian cortex.

Q: The "neural" metrics focus on cellular connectivity and activity – can one propose to study smaller-scale phenomena like nonlinear dendritic processing or larger-scale phenomena like local-field potentials?

A: Yes. The metrics specify *minimum* requirements. In many cases, we expect performers will exceed these requirements and "zoom" in or out (in space, time, and other relevant dimensions) on structures and dynamics of particular interest.

Preemptive FAQ (Part 4)

Q: Will MICrONS fund development of neuromorphic hardware?

A: Probably not, unless hardware development is required to support computational neural modeling or machine learning activities. Note that efficient execution of models and algorithms is not the focus of this program.

Q: Will MICrONS fund development of new experimental or imaging techniques?

A: Maybe, if a case can be made that the new techniques will (a) be ready for use by the time they are required in the program schedule, and (b) enable capabilities that significantly increase the probability of achieving MICrONS program goals. MICrONS does not intend to fund basic research in tool development.

Q: Can one propose to a single technical area?

A: No. The technical areas all interact – directed and integrated research efforts are likely required to achieve our goals with the limited time and money available.

Q: Can one be on multiple proposals?

A: Yes.

Q: Are National Laboratories, FFRDCs, UARCs, or other similar types of organizations that have a special relationship with the Government eligible to participate on a team?

A: No, but personnel at these organizations with joint appointments at universities or other eligible institutions are welcomed to submit a conflict waiver for consideration on a case-by-case basis.

Q: Are foreign (non-U.S.) individuals or organizations eligible to participate on a team?

A: Yes.